

Stochastic Inversion of Electrical Resistivity Changes Using a Markov Chain, Monte Carlo Approach

Abelardo L. Ramirez
John J. Nitao
William G. Hanley
Roger D. Aines
Ronald E. Glaser
Sailes K. Sengupta
Kathleen M. Dyer
Tracy L. Hickling
William D. Daily

This article was submitted to the Journal of Geophysical Research

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

August 2003

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

Stochastic inversion of Electrical Resistivity Changes Using a Markov Chain, Monte Carlo Approach

A.L. Ramirez, J. J. Nitao, W. G. Hanley, R. Aines, R. E. Glaser, S. K. Sengupta, K. M. Dyer, T. L. Hickling, W. D. Daily

Lawrence Livermore National Laboratory, Livermore CA, 94550
August, 2003

Abstract

We describe a stochastic inversion method for mapping subsurface regions where the electrical resistivity is changing. The method combines prior information, electrical resistance data, other relevant data and forward finite difference models to produce subsurface resistivity models that are most consistent with all available data. Bayesian inference and a Metropolis simulation algorithm form the basis for this approach. The approach enables the estimation of distributions of both individual parameters such as center of mass and groups of parameters such as resistivity change images. Attractive features of this approach include its ability to: 1) provide quantitative measures of the uncertainty of a generated estimate, 2) seamlessly integrate disparate data such as electrical resistance measurements and liquid volume measurements, 3) effectively invert complex nonlinear forward models without appealing to unrealistic simplifying assumptions, 4) function effectively when exposed to degraded conditions including: noisy data, incomplete data sets and model misspecification and, 5) allow alternative model estimates to be identified, compared and ranked. The proposed method is computationally expensive, requiring the use of large computer clusters to make the approach practical. A series of physical model test cases have been performed to validate the approach. Field results using data collected during the infiltration of a salt-water tracer are also discussed. Methods that assess MCMC convergence and summarize interesting features of the posterior distribution are also introduced. The stochastic inversions presented suggest that zones of resistivity change can be successfully mapped with this approach. The stochastic tomographs accurately identify the most probable location, shape, and volume of the changing region, and the most likely resistivity change.

Introduction

A fundamental earth sciences problem is to determine the properties of an object that we cannot directly observe. A variety of geophysical tomography techniques have been developed to provide detailed subsurface information. One such technique, electrical resistance tomography (ERT), is a relatively recent geophysical imaging technique that provides 2-D and 3-D images of resistivity that are consistent with measurements made on an array of electrodes. With the increasing availability of computer controlled multi-electrode instruments and robust data inversion tools, ERT is becoming widely available.

The value of ERT for monitoring dynamic subsurface processes has promoted new applications in a wide range of environments (e.g. see review article by Daily et al., 2001). Specifically, ERT has been applied to a variety of environmental and engineering scales and processes such as imaging of flow and transport in porous media and fractured rock at small and large scales; unexploded ordnance detection; municipal landfill site characterization; and monitoring the progress of in-situ remediation processes.

The goal of any ERT inversion method is to calculate the subsurface distribution of electrical resistivity from a large number of resistance measurements made from electrodes. A deterministic inversion procedure searches for a single model (i.e., a spatially varying distribution of resistivity) that gives an ‘acceptable’ fit to the data and satisfies any other prescribed constraints. A common solution to this inverse problem minimizes an objective function consisting of a regularized, weighted least squares formulation. Typically, the search is conducted using iterative, gradient-based methods. (e.g., Ellis and Oldenberg, 1994, LaBrecque et al., 1996).

The solution to the classical (deterministic) ERT inversion problem yields one model of resistivity structure that then can be used to predict system behavior. The inversion problem is typically complicated by a non-linear relationship between data and the inverted parameters, state-space dimensionality, under/over determined systems, noisy and dependent data, etc. Hence, an exact inversion is rarely possible. In fact, without appealing to unrealistic simplifying assumptions, the severity of these issues often makes classical optimization algorithms ineffective at estimating (i.e., inverting) those system parameters that most consistently correspond to the observed data. Moreover, conventional solutions provide little insight into the degree of uncertainty associated with the inversion result.

Previous Work— Stochastic Methods:

One alternative to the classical ERT implementation is to use stochastic methods that search for electrical resistivity models that best fit the collected electrical resistance measurements. The literature describes a variety of stochastic methods used in geophysical applications. Zhang *et al.* (1995) suggest an inversion method that seeks to maximize the *a posteriori* probability density function of model parameters. Although the method itself is very general, implementations of this method for ERT rely on several restrictive assumptions including: a) the resistivity parameters have known covariance, which depend only on the distance between the physical locations of the parameters, and b) the data errors and the parameters have normal (i.e., Gaussian) distributions. With these and other assumptions, maximizing the *a posteriori* probability density function is equivalent to minimizing the objective function in the classical inverse approach. Yang and Labrecque, 1999, proposed an alternate solution that allows a more reasonable estimate of the parameter covariance matrix by avoiding the need to directly invert it.

Mosegaard and Tarantola(1995) and Mosegaard and Sambridge(2002), describe a second statistical approach to inversion problems. In this approach, the inverse problem is

reformulated as a Bayesian inference problem based upon an estimate of the *a posteriori* probability distribution generated via the Metropolis algorithm. The underlying posterior distributions combine available prior information and physical models with new information obtained through the measurement of some observable parameters. The technique utilizes the Markov Chain Monte Carlo (MCMC) method and a well-known importance sampling method known as the Metropolis algorithm. A key characteristic of this technique is that solutions are sampled at a rate proportional to their posterior probabilities. Hence, models consistent with *a priori* information as well as observations are sampled most often, while models that are incompatible with either prior information and/or observations are rarely sampled. It allows the use of complex *a priori* information, and data with an arbitrary noise distribution.

Mosegaard and Tarantola(1995) applied this algorithm to a highly non-unique, linear inverse problem to show the method's ability to extract information from noisy data. They also describe an approach that allows the simultaneous use of disparate types of data; in geophysics we may have different data sets such as electrical, seismic, and gravity data. The authors use a "cascade rule" approach to process these different data sets in an integrated fashion assuming the associated measurement errors are independent. When the solution of the forward modeling is inexpensive for one (or more) of the data sets, the "cascade rule" may render the algorithm much more efficient than the using the Metropolis rule simultaneously on the entire data set

We will refer to the approach using the cascade rule as SIMCS (statistical inversion and Monte Carlo sampling). SIMCS can be used to solve highly non-linear, discrete inverse problems such as typical ERT inverse problems. For these problems, we generally have incomplete knowledge of the relationship between data and model parameters. SIMCS provides the means to perform a thorough resolution and uncertainty analysis. It also efficiently samples all local maxima of the posterior distribution, provided a sufficient number of iterations are performed.

Kaipio et al., 2000, describe the application of the SIMCS approach to medical imaging problems using electrical impedance tomography (EIT, for the purpose of this paper, ERT and EIT are synonymous). The authors point out that the problem is to estimate the posterior distribution of the unknown model resistivities conditioned on the measured transfer resistances. All the variables included in the model are characterized as random variables. The randomness is related to our perceived uncertainty of the true model values, and this uncertainty is expressed in terms of probability distributions. Kaipio et al considered a variety of priors including a minimum total variation prior, a second order smoothness prior and an "impulse" prior that penalizes the L^1 norm of the resistivity. They indicate that the posterior probability distribution measures our uncertainty based on the resistance measurements and the prior.

Andersen et al. (2001a) describe a SIMCS approach for the Bayesian inversion of geoelectrical resistivity data. They used random, polygonal models to represent the layered composition of the earth, and demonstrate the performance of the method using field data. They also describe the use of the multi-grid, Metropolis-coupled, MCMC

approach to improve the efficiency of the method. In this approach, several Markov chains are launched within dissimilar state spaces, and then randomly swap states between the chains. They suggest that the multi-grid approach converges faster than the reversible jump MCMC algorithm.

Andersen et al. (2001b) describe another SIMCS application aimed at the detection of cracks in electrically conductive media. Their approach assumes that the cracks are linear, non-intersecting and perfectly insulating. Using synthetic data, they demonstrate a Metropolis-Hastings updating scheme that assumes that the number of cracks is known *a priori*. They also considered a reversible jump MCMC approach that allows updating of the number of cracks, and a simulated tempering method that increased the computational efficiency of the overall process. They conclude that their approach is promising and robust, but also computationally intensive.

Our Approach -- Stochastic Engine

Here we describe a type of SIMCS algorithm that incorporates resistance measurements, numerical forward simulators of subsurface electrical resistivity, and *a priori* knowledge to provide distributions of resistivity change that are likely to be present in a subsurface environment. The approach is based upon a statistical inference paradigm called the Stochastic Engine (SE) (Aines et al., 2002, Glaser et al., 2002, Newmark et al., 2001). This methodology combines disparate types of observational data such as cross-borehole electrical resistance measurements, and liquid volume data to produce a consolidated body of knowledge indicating those subsurface plume configurations and system parameter values which are most consistent with the available data and forward models (i.e., system identification).

As with all SIMCS approaches, Bayesian inference and a Metropolis-type simulation algorithm form the basis for our approach. There are two major components to the approach:

- i) A base representation specifying the form that the resistivity configurations (i.e., states) of the system may assume. This representation provides: i) a framework for incorporating up front problem specific information and restrictions, ii) the basis for generating the forward model predictions conditioned on an assumed system state, and iii) a specification of a general element in the support of the estimated posterior distribution.
- ii) A Markov Chain Monte Carlo (MCMC) simulation algorithm that generates samples according to the unknown posterior distribution. This is accomplished by first proposing samples from a distribution of possible resistivity configurations that are consistent with our prior data. Then, we use a Metropolis-type decision algorithm to accept or reject the proposed states according to their consistency with the observed data. After a sufficient warm-up period (called burn-in), the accepted states constitute a sample from the posterior distribution.

There are two key differences between the SE approach and the resistivity inversion approaches described earlier. In the SE approach, we make use of multiple disparate data

types whereas these other approaches typically use a single data type (resistance data). Furthermore, the goal of the other approaches is to find a single resistivity model that is consistent with the data, whereas the goal of the SE inverse process is to find alternative models that are consistent with all available data, and to rank them according to their probabilities. Details of the SE inverse approach are presented next.

Base Representation - Subsurface Plumes

The representation of the state space S (i.e. defining the individual resistivity configurations and their neighborhood structure) is critical to the overall effectiveness of the method. In those cases where traversal of the state space is expensive and/or time intensive (e.g., computationally intensive forward models such as those used for 3D ERT), the state representation must be designed so that the number of degrees of freedom in the problem is kept within reasonable limits. The resistivity perturbation model used to specify neighboring states is very efficient in this respect. We use a stochastic model that generates neighboring configurations based upon an identified spatial distribution of resistivity change. The properties of subsurface plumes often depend on lithology or facies and because of this, we have chosen a categorical simulation approach. The main advantages of this approach are: (1) the range of possible resistivity magnitude is [binned into a collection of discrete categories](#) thereby constraining the size of S , and (2) geologic and hydrologic insight on the spatial characteristics of the subsurface plume can be exploited.

Liquid leaks from tanks and subsurface fluid injection (e.g., steam, CO₂, or water floods) can produce regions where the resistivity changes as the released liquid penetrates the rock or soil mass. When time-lapse electrical resistance measurements are made before and during a release sequence, changes in the electrical resistance data can be used to map the subsurface regions penetrated by the liquid. In this case, the goal of the SE simulation process is to determine the most probable shape, location, resistivity, and volume of the region penetrated by the liquid.

Our base representation takes advantage of prior information about subsurface plume characteristics. It assumes that the system of interest consists of a zone of changing electrical resistivity embedded within an otherwise homogeneous volume. We further assume that the changing volume can be described with 2 – 10 contiguous sub-volumes. The range of possible resistivity changes is assumed known a-priori. While the sub-volumes have to be contiguous, they can overlap and have varying size, shape and resistivity properties thereby allowing the modeling of plumes with complicated resistivity structure.

For some subsurface injection applications, other data such as the volume of the injected liquid, temperature, pore water electrical conductivity is routinely collected. One advantage of the base representation framework is that these other data can be used to constrain characteristics of the representation for the plume region. For example, in situations where the released liquid volume is known, the volume of the proposed region is forced to be consistent with the released liquid volume. This is accomplished by using a suitable petrophysical model (e.g., Archie's equation or the Waxman-Smiths equation

(Hearst et al. 2000) that relate electrical resistivity to other parameters of interest such as liquid saturation of the pore space.

Markov Chain Monte Carlo

The Markov Chain Monte Carlo method (MCMC) provides a flexible framework that can be adapted to perform a variety of analyses and inference tasks. It uses a Markov chain state/transition structure to control the sampling process. MCMC techniques enable simulation from a targeted distribution by embedding it as a stationary distribution of an ergodic Markov chain and simulating the chain until it approaches equilibrium. For Bayesian analysis, we are able to adapt the approach to simulate and estimate posterior distributions that embody our available prior information and newly acquired observational data. MCMC algorithms can assume a variety of forms with the most useful to us being the Metropolis framework.

The SE approach is a derivative of the Metropolis algorithm (Metropolis et. al., 1953) as described by Mosegaard and Tarantola (1995). This particular MCMC algorithm has demonstrated significant potential in solving inverse problems involving complex physical systems and supports several key enhancements necessary to mitigate the combinatorial demands underlying the MCMC methodology. Within our stochastic framework, the solution to an inverse problem is an estimate of the posterior probability distribution defined on the corresponding space S of possible solutions (otherwise referred to as states). For any potential solution $s_0 \in S$, the SE will provide an estimate of the probability and confidence that state s_0 is indeed the true solution to the given system. This allows future analysis to focus upon the most likely explanations of system behavior – thereby improving both the efficiency and effectiveness of follow-on efforts. Moreover, results generated from the SE (i.e., the estimated posterior distribution, predictions, hypothesis testing, etc.) may be incrementally updated as more data becomes available.

The inverse problem under consideration may be described as follows. Let D and M denote the data space and model space respectively, and suppose that there exists a mapping G such that:

$$\underline{d} = G(\underline{m}) \quad (2)$$

where $\underline{m} \in M$ is a parameter vector describing the state of the system of interest and $\underline{d} \in D$ is a vector of measurements taken on that system. The inverse problem occurs when a vector of data values is observed, say \underline{d}_0 , and we want to determine the value of the parameter vector \underline{m}_0 that gave rise to \underline{d}_0 . In geophysical ERT applications, this problem is substantially under-constrained and ill-posed. In such cases, the search for a deterministic solution for \underline{d}_0 that is unique and possesses a high degree of confidence is virtually impossible, and hence a probabilistic solution is likely to be superior to any classical deterministic optimization approach.

The stochastic approach employs a variation of the Metropolis algorithm to generate a sequence of samples from M at rates proportional to their posterior probabilities. These

samples enable the estimation of the posterior distribution $F_{M|D}$ using the sample frequencies corresponding to the sampled models $\underline{m} \in M$. By design, the models generated most frequently are most consistent with both our prior information on M and the observations being processed. Assuming mild regularity conditions, the estimated posterior distribution $\hat{F}_{M|D}$ converges strongly to the true underlying distribution $F_{M|D}$. This estimate is the basis for all subsequent analysis and inference including estimation, prediction, confidence assessment and risk/reward trade-off analysis.

Given that the information used to drive the simulation is taken from two distinct sources (prior information and observational data) the sampling process can be viewed as consisting of two separate components. The first component generates samples according to an identified prior distribution $\rho(\underline{m})$ defined over the model space M . These samples are called proposal states and constitute possible solutions to the inverse problem, (refer to the right hand side of Figure 1, base representation box). Specifically, this sampling process is manifested as a Markov chain, \mathbf{Q} , with one-step transition probabilities between states $\underline{m} \in M$ designed to produce a long-run stationary distribution equal to the prior $\rho(\underline{m})$. In statistical terms, \mathbf{Q} samples $\rho(\underline{m})$.

The second component takes the form of a decision process that either accepts or rejects the proposal sample generated from the a-priori Markov chain (shaded area labeled “stage 1” in Figure 1). For each visited state, forward simulators are used to predict values of measurable parameters such as electrical resistance. These predictions are then compared to corresponding measurements to determine the likelihood $L(\underline{m})$ that the given state $\underline{m} \in M$ produced the observed data. An accept/reject decision based upon this likelihood is used to modify the prior sampling process. The result is a new Markov chain, \mathbf{R} , which samples the posterior distribution, $\rho(\underline{m})$.

Formally, Bayes rule relates the prior and posterior distributions as follows:

$$\rho(\underline{m}) = k \rho(\underline{m}) L(\underline{m}) \quad (3)$$

The likelihood $L(\underline{m})$ is a measure of the degree of fit between the data predicted assuming the model \underline{m} and the observed data, and k is a normalization constant. For this study, we assumed a likelihood function of the form:

$$L(\underline{m}) = k \exp - \left| \left(d_{pred,i} - d_{0,i} \right)^n / \sigma_i^n \right| \quad (4)$$

where $d_{pred,i}$ is the predicted data for a given model \underline{m} , $d_{0,i}$ is the vector containing the observed measurements, σ_i is the estimated data uncertainty, and $n = 1$. For the results described below, we assumed that $n = 2$. Note that Eq. 4 assumes that the estimated data errors are uncorrelated.

The decision to accept or reject a proposed state is made on the basis of likelihood comparisons as indicated by the “stage 1” box, right hand side of Figure 1. Suppose that the current state of the Markov chain is \underline{m}_i and that a randomized rule based upon the one-step transition probabilities propose a move to state \underline{m}_j . If these transitions were always accepted, then the simulation would be sampling from the prior distribution. But, instead suppose that the proposal transition is only accepted according to the following rules:

- 1) For both the current and proposal states \underline{m}_i and \underline{m}_j , compute the respective likelihoods $L(\underline{m}_i)$ and $L(\underline{m}_j)$ that these models produced the observed data.
- 2) If $L(\underline{m}_j) \geq L(\underline{m}_i)$, then accept the proposed transition with probability 1 and move to state \underline{m}_j . (Note: The algorithm always accepts the transition when the new state provides a better explanation of the data than the current state.)
- 3) If $L(\underline{m}_j) < L(\underline{m}_i)$, then use a randomized decision rule and accept the proposed transition with probability $L(\underline{m}_j)/L(\underline{m}_i) < 1$ and move to state \underline{m}_j . Otherwise, transition back to state \underline{m}_i . (Note: By allowing the random walk to transition to a less likely state, the process can move out of a local extrema).

The samples generated through this three-step process will have a limiting distribution that is proportional to the desired posterior distribution $(\underline{m}|\underline{d}_0)$. That is, the search tends to hover in regions of space M containing states with greater prior propensities and higher likelihoods.

Since we have observational data of differing types (e.g., ERT and flow data), our approach takes advantage of the “cascade” rule (Mosegaard and Tarantola, 1995) In these cases, the errors in prediction are often independent and hence the total likelihood expression factors into distinct terms – one for each data type. Hence, for the above example the total likelihood expression factors as follows:

$$L_{total}(\underline{m}) = L_{ERT}(\underline{m}) \cdot L_{flow}(\underline{m}) \quad (5)$$

This probabilistic structure can be leveraged to streamline the transition process employed by the SE. Mosegaard and Tarantola (1995) indicate that performing a single Metropolis transition step (step 3 above) that uses the entire likelihood expression $L_{total}(\underline{m})$ is equivalent to performing a sequence of Metropolis steps – one for each term in the above expression. Once a model is proposed by the prior distribution, the forward model is solved initially for the first data type alone (step 1 as illustrated by the stage 1 box in Figure 1). At this juncture, the proposed model may be rejected or accepted (steps 2 and 3 above). If the decision is to reject the proposal, then the forward models in stage 2 (Figure 1) are not executed. The Markov chain returns to the current state, a new state is proposed and the decision process begins anew with the first data type. If the decision at this stage is to accept the proposal, the next data set is considered in stage 2, its corresponding forward model is run and a decision to accept or reject is made based upon its likelihood. This continues through all of the different types of data until the proposal

either is accepted at all stages or is rejected at one stage and starts over at the beginning of the sequence. Note that proposals accepted by both stages are ones that are most likely to be consistent with prior data and stages 1 and 2 observational data. Also, this approach can easily be extended beyond two stages to incorporate additional types of data. The results we show in the next section made use of only a single stage.

This staging approach allows the processing order of the likelihood terms to be arranged according to increasing runtime complexity of the corresponding forward model since the order of the stages does not affect the posterior estimate. However, when forward models in the various stages require vastly different computation times, it is wise to put the faster model in the first stage because generally a higher number of models are evaluated in the initial stage. Thus, stage order does affect overall computational time of the simulation.

The right side of Figure 1 illustrates how the MCMC search locates regions of high likelihood within space S . The bottom landscape represents space S ; the states within this space are all consistent with prior information. Also shown is an example state that is one of the many potential states, $s_0 \in S$. The height of the hills identifies regions of S where states that are consistent with prior information are located. Note that some of the peaks grow taller (and others shorten) as the Markov chains(s) tend to hover in regions that are consistent with data analyzed in stages 1 and 2. The taller peaks identify regions where states are most consistent with the data (have a higher likelihood); states that are inconsistent with the data occupy the flatlands. We will refer to the peaks as modes in subsequent sections. The landscape associated with stage 2 contains samples of the posterior distribution $(\underline{m}|\underline{d}_0)$ that will be analyzed to draw conclusions about the system under study. The presence of multiple modes in the posterior distribution indicates that there are alternative models that explain the data. The likelihoods of the different modes (i.e., peak heights) provide a way with which to rank the alternative models; the highest likelihood mode contains models that best explain all available data. The width of the peaks indicates the degree of state variability within each mode.

One major advantage of our stochastic approach is that knowledge of the posterior distribution allows the uncertainty of the generated estimates to be quantified. This provides the basis for: i) the objective assessment of competing hypotheses when the available information isn't sufficient to definitively identify the system state, and ii) the propagation of uncertainty in modeling results through to follow-on predictions. Sources of uncertainty such as measurement error, contradictory data, lack of sensitivity or resolution, incomplete surveys, and non-unique relationship between measurements and inverted parameters are explicitly considered by this approach. Problems with many secondary extrema, a non-unique inverse, and/or contradictory or sparse data can be inverted using the SE. Even when conventional inversion and analysis methods are able to address complex problems, they provide only a single "best" answer, throwing away much information and precluding other likely possibilities. Alternatively, the SE allows continuous integration of new data into the analysis, improving understanding and reducing uncertainty. In addition, the methodology is generally applicable to ill-posed, highly non-linear, poorly constrained, multi-dimensional problems due to its forward processing scheme (i.e., the simulators are all forward models).

There are a variety of key issues that must be addressed during the implementation of this methodology. The most fundamental concern is that the Markov chain must be designed so that a limiting stationary distribution actually exists and the process converges. For this to happen, the transition probabilities must be defined so that the process is aperiodic (state transitions are not cyclical, state sequences do not repeat) and irreducible (it is possible to move from any given state to any other). Our approach also assumes that the transition probabilities are not known a-priori. This assumption requires that:

$$p(\text{state}_x \rightarrow \text{state}_{x+1}) = p(\text{state}_{x+1} \rightarrow \text{state}_x) \quad (6)$$

That is, the probability of moving from one state to another is the same as the probability of moving in the reverse direction.

Once we are assured of convergence of the simulation, the critical issue becomes its rate of convergence. Key factors affecting the convergence rate include: (i) the strength and quality of the prior distribution $p(m)$, (ii) the representation, resolution and dimensionality of the state space M , and (iii) the structure of the likelihood surface. The convergence rate of the process and the forward model runtimes effectively control the efficiency and runtime of the overall SE simulation.

Base representation resolution is critical; if the resolution is too fine or involves a high dimensional state vector, the convergence may be slowed beyond practical limits. The step size chosen (i.e. the resolution of the state space) controls the neighborhood size around a given state. If too small, movement through the state space will be slow, and it will take longer to move past local extrema in state space M . When the step size is too large, increased rejection ratios (number of states rejected/total number of states evaluated) are likely thereby slowing the convergence rate. By trial and error, we discovered that a randomized step size that sometimes took smaller or larger steps provided a reasonable solution to this dilemma.

Closely related to the step size is the structure of the likelihood surface being traversed. When the surface is exceptionally steep, the process will be slowed due to high rejection rates. This situation occurs when the process attempts to move off a steep peak in likelihood space and ends up proposing states having much smaller relative likelihoods that are almost always rejected. This slow mixing process (i.e., how well the process moves through the state space) can be mitigated in several ways including the use of multi-resolution or the Langevin diffusion methods. Finally, the choice of the prior distribution may also significantly impact convergence it controls the proposal of candidate states. In general, the closer the prior is to the posterior, the faster the process converges.

Posterior Analysis

The posterior distributions estimated via the MCMC provide the basis for inference by performing a variety of subsequent analyses including parameter estimation, confidence

assessment and risk analysis. Posterior analyses methods will be discussed in detail subsequent sections.

Resistivity Change

Electrical resistivity changes are typically created by hydraulic, chemical and thermal changes in the subsurface. Thus, monitoring natural or artificially induced changes in resistivity can often provide valuable information about flow and transport processes, including the effectiveness of remediation strategies for clean up of contaminated sites. Applications such as monitoring the progress of water and liquid contaminant movement, monitoring of cleanup processes such as air sparging, and steam injection, and the detection of leaks from tanks containing hazardous liquids have been reported in the literature (Binley et al., 1996, 2001, Daily et al., 1987, 1992, Kemna et al., 2000, LaBrecque et al., 1996, Ramirez et al., 1993, 1996, Slater, 1997). In all these case, electrical resistance measurements are made before and after the start of a process that changes the subsurface electrical resistivity distribution. Tomographs of electrical resistivity change can be constructed in variety of ways such as: a) by inverting the “before” and “after” data to produce tomographs of electrical resistivity, and then subtracting the two tomographs, or b) by calculating differences in the raw transfer resistances and inverting the differenced data. LaBrecque and Yang (2000) developed an approach based on the latter strategy that uses a modified data vector formed from differences in measurements. Their approach not only conditions parameter estimates on some background state but also significantly reduces computational effort in the inversion.

An alternative approach related to strategy “b” utilizes a ratio of two impedance datasets in the inversion. In this method, a new data vector, \mathbf{d}_r^* , is formed from:

$$\mathbf{d}_r^* = \frac{\mathbf{d}_t^*}{\mathbf{d}_0^*} \mathbf{f}^*(\sigma_{\text{hom}}^*), \quad (1)$$

where \mathbf{d}_0^* is the data vector at some reference state, \mathbf{d}_t^* is the data vector at some time t and σ_{hom}^* is an arbitrarily chosen homogenous conductivity.

Inversion of the new dataset \mathbf{d}_r^* in the normal manner then results in an image that will reveal changes relative to the reference value σ_{hom}^* . Most applications of this approach have been in DC resistivity problems where only impedance magnitude (i.e., resistance) is used (see Daily *et al.*, 1992; Ramirez and Daily, 2001); however, it may also be applied to the general complex resistivity problem. We will use this resistance ratio-approach to produce the results described later in the paper.

Results and Discussion:

Physical Model Results

We have used physical models with a state space structure that is readily diagnosed in order to develop a better understanding of the Stochastic Engine approach. The model used consisted of various targets immersed in a tank filled with water. The model included 4 vertical electrode arrays, each having 15 electrodes (refer to Figure 2). The arrays were submerged in water and a variety of solid acrylic and porous (sand-lead mixture encased in a nylon mesh) targets having different electrical resistivities were inserted at various locations between the electrodes. The water resistivity was 16 ohm-m, the sand-lead mixture had a resistivity of about 40 ohm-m, and the acrylic resistivity was approximately 10^{10} ohm-m.

Issues such as the accuracy of the inverted location, shape and change magnitude of the inversions were evaluated. For these inversions we assumed that the target could be sufficiently described by two contiguous or overlapping parallelepipeds. Their shape, location and contrast were allowed to vary. Uncertainty arises from the inherent errors (measurement and modeling) and the non-unique relationship between inverted parameters and measurements.

Figure 3 compares the stochastic inversion results for the sand-lead target and the equivalent deterministic inversion. The top left image is a vertical section of the actual target showing its shape and location. Given resistivity contrasts in typical geophysical applications, the target has a low-contrast (resistivity only 2.5 times as high as the surrounding water). We assumed that the resistivity change (as indicated by a ratio) was discrete and between 0.5 and 16.0 (0.5, 1.0, 2.0, ..., 16.0); 1.0 indicates that the target's resistivity equals that of the surrounding water, and values above 1.0 indicate that the target's resistivity is larger than the surrounding water. The inverse process searched for the most likely location, size, shape and contrast of the target.

The relationship between the base representation and the posterior distribution is straightforward for this problem. We summarized sample information by calculating voxel-wise means of all the images generated during the simulation of the posterior distribution. The Fig. 3 top middle and right frames show the voxel-wise average resistivity ratio for all the sample images. These graphs show similar shape and location as the actual target. The resistivity ratio magnitude lies in the range of 2.0 – 2.9, while the actual target value is 2.5. The top right image shows an iso-surface representation.

The lower row of images in Figure 3 shows a deterministic tomograph calculated by the inversion algorithm described by LaBrecque et al, 1999. The bottom left image displays the results using the same color bar range as the stochastic image in the middle of the top row. The deterministic inversion appears at approximately the same location as the stochastic result but shows a resistivity ratio (about 1.2) that is much lower than both the target and the stochastic image. The bottom middle and right images show the deterministic result using a much narrower color bar range so that image details become easier to observe. The bottom middle image shows that the lateral and vertical extent of the deterministic anomaly are somewhat exaggerated.

The larger size of the deterministic result is caused by the approach seeking to find a unique, robust solution to the inverse problem. Direct linear inversion is not possible because it only applies to linear problems and the ERT problem is generally non-linear and underdetermined. The typical ERT problem is ill-conditioned, and therefore requires regularization in order to stabilize the solution and obtain meaningful results (Park and Van, 1991, Shima, 1992, Ellis and Oldenburg, 1994, Sasaki, 1994, LaBrecque et al, 1996, Morelli and LaBrecque, 1996). The most common regularization used constrains the parameter search by penalizing models with large “roughness” (the inverse of smoothness). The deterministic inverse algorithm uses a roughness-constrained, least-squares, iterative approach that searches for a single model that fits the data within a specified tolerance. Using this approach it is possible to obtain meaningful “smooth” parameter models even when the problem is underdetermined.

A disadvantage of this methodology is that roughness-based regularization distorts the resistivity values between adjacent elements thereby producing models that have reduced contrast and exaggerated extent. This is the reason for the somewhat distorted anomaly in the deterministic result shown in Figure 3. The size of the solution space is much larger for the deterministic approach because it includes many more possible target configurations such as discontinuous zones, infinite number of resistivity contrasts etc. We note that it is possible to incorporate prior information in deterministic inversions (e.g., Rodgers, 1976, 1990) but we did not consider it here.

Next, we consider a very high contrast (approx. 10^9), acrylic target. The inversion results are presented in Figure 4. We assumed that the possible resistivity contrast range was somewhere in a range of discrete steps (10^0 , 10^4 , 10^5 , 10^6 , 10^7 , 10^8). The inverse process searched for the most likely location, size, shape and contrast of the target. The top left frame shows the actual target located on the left side and just below the center of the 3D block. The lower left and lower right frames show two different views of the stochastic inversion results. We elected to show voxel-wise location probability instead of the average resistivity ratio because the latter gives very misleading results when the contrast between target and surrounding water is very large. Suppose that our posterior distribution produced 10,000 samples, and that at a given voxel, 9999 samples have a value of 10^0 and only one sample has a value of 10^7 . For this case, the average resistivity ratio is about 10^3 even though practically all voxels have a value of 10^0 . Thus, instead of a voxel-wise resistivity ratio, we chose to use voxel-wise location probability. The latter is calculated on a per voxel basis, by counting the number of posterior samples that contain a portion of the anomaly and dividing by the total number of posterior samples.

The results in Figure 4 show that the location and size of the inverted anomaly are similar to the actual target. It is interesting to note that our approach produces anomalies with fairly sharp boundaries even when the target to background contrast is many orders of magnitude. The right frame in the figure shows the inversion results for a “blank” target; i.e., two measurement surveys were collected a few hours apart with only water present. Any changes in the data are the result of measurement error. The “blank” result shows how these errors propagate through the stochastic inversion process. Note that while the

voxel-wise probabilities are not equal to 0.0, they are much smaller than those obtained when the acrylic target is present.

The models considered in Figures 3 and 4 are analogous to geophysical cases where one is interested in detecting subsurface changes in resistivity caused by various processes such as hazardous liquids leaking from tanks, and steam or CO₂ injection. We next look at an application of the SE method using field data.

Field results - Tank Leak Detection

Large volumes of hazardous liquids are stored worldwide in surface and underground tanks. Frequently these tanks are found to leak, thereby resulting in not only a loss of stored inventory but, more importantly, contamination to soil and groundwater. There are two methods of detecting leakage from tanks: monitoring the liquid level of waste in the tanks and monitoring the soil surrounding the tanks for leaks. When the spillage of the liquids changes the electrical resistivity of the soil or rock in a measurable way, the latter monitoring strategy maps the electrical resistivity around and below a tank over time.

Electrical resistance data was collected during a field experiment that simulated leakage from a large metallic tank. For testing, an electrical equivalent (saline solution) was used instead of the real contaminant to preserve the environmental quality of the test site. The test site used for this work is part of the 200 East Area in the Hanford Site, located near Richland, Washington (additional details about the test site and testing are provided by Barnett et al., 2002). The near surface sediments at the test site consist primarily of fine to coarse-grained sand displaying plane lamination and bedding. The field experiments were performed under a 15.2 m diameter steel tank mockup. Figure 5 shows the layout at the leak detection experiment site. This empty steel tank contained several built-in spill points (one of which is shown). The bottom of the tank is located 1.5 m below ground surface. Sixteen boreholes with eight electrodes in each surrounded the tank. The electrodes were spaced every 1.52 m between the surface and 10.7 m depth. The water table was located well below the deepest electrode location. The diametrical distance between boreholes was 20.7 m.

Hypersaline solution was released from a point near the tank's center (refer to Fig. 5) over a 52 hour period. The release rate averaged about 43 liters/hour. The liquid consisted of a sodium-thiosulfate solution with a conductivity of about 5 S/m (36 wt%). We note that this hypersaline fluid has similar electrical conductivity and density to the real Hanford tank liquids. ERT surveys were made before, during and after the brine release using a dipole-dipole approach.

Figure 6 shows the results of a stochastic inversion (left panel of images) and the corresponding deterministic inversion (right panel of images). The results represent the inverted resistivity change caused by the release of approximately 2160 liters of brine solution. The dimensions of the 3D block are 28 m (along each horizontal axis), and 13 m (height). The stochastic results show the pixel-wise average resistivity ratio, calculated using approximately 1200 posterior samples. The deterministic results were calculated using the same gradient search method described in the prior section. The slices shown in

the top row of images are oriented North-South. The bottom row of images shows iso-surface views, where the white bar over the color bar indicates the range of transparencies used to render the iso-surfaces.

Both methods show that the resistivity decreased (ratios below 1.0 represented by the blue and green colors, top row of images) directly below the release point, and the zone of change extends more to the East of the release point than to the West. There are also differences between the images. The stochastic image shows much stronger resistivity changes reaching values of 0.2 (i.e., the invaded soil's has 1/5 the resistivity of the pre-release value). Also, the changing region in the stochastic image shows a smaller volume. The deterministic image shows a substantially weaker change of 0.8. The stochastic anomaly is largely constrained to immediately below the release points whereas the deterministic anomaly shows near surface changes near the block corners and away from the release point.

A drill-back program that would have independently mapped the plume's extent was outside the scope of this work. Instead, we are forced to use circumstantial evidence to evaluate the results. Ward and Gee (2001), and Gee and Ward (2001) describe the results of field tests, laboratory tests and numerical simulations of the movement of hypersaline solutions through the vadose zone. Their tests were conducted at a site located a few hundred yards from the Mock Tank leak test facility where our results were obtained; both sites are known to have similar geology. Ward and Gee suggest that hypersaline fluids move through the vadose zone along finger-like preferential flow paths due to the much higher density of the hypersaline solution. Gee and Ward point out that hypersaline plumes tend to be more compact, move deeper and show less lateral spreading than low ionic strength solutions.

We propose that the stochastic result in Figure 6 is a more realistic representation of the plume than the deterministic result due to the following reasons.

- 1) The resistivity change in the stochastic anomaly is much stronger; we note that the released hypersaline fluid is about 100 to 1000 times more electrically conductive than the native pore water; the exact figure is difficult to ascertain due to the effects of previous conductive tracer releases at site. However, we believe that the stronger changes in the stochastic image (values near 0.2) are more likely to be representative of the plume's properties because of the high conductivity of the released fluid.
- 2) The stochastic anomaly is more compact, and finger-like, thereby suggesting preferential flow path behavior consistent with the observations of Ward and Gee (2001) and Gee and Ward (2001). The deterministic anomaly shows more lateral spreading due to the regularization approach used.
- 3) The deterministic anomaly shows unrealistic features such as those located near the top of the image block, away from the release point; in contrast the stochastic image shows the bulk of the changes concentrated directly below the release point.

Most tank managers want answers to the following. a) Is there a leak? b) How much liquid has been released? c) What are the location, volume, and extent of the

contaminated zone? We attempt to answer some of these questions using the sequence of inversions shown in Figure 7. The inversions represent changing subsurface resistivity as the volume of released fluid increases from 340 to 2160 liters. Along a given column, released volume increases from the top row to the bottom row of images. The left column images shows voxel-wise location probability. We believe that the location probability images help us determine if there is a leak. The acrylic and “blank” target results (shown in Fig. 4) are analogous to “leak” and “no leak” conditions in the tank leak application; note that the “blank” results in low probabilities (below 0.1) that extend over a broad region of the 3D block. When the acrylic target is introduced, the probability increases to about 0.2 in a localized region and drops to near 0 on the rest of the block. The tank leak results (Figure 7) show similar qualitative trends. At low volumes: the location probabilities average about 0.10 and extend over a broad, ellipsoid shaped region. At higher volumes the location probability increases to values near 1.0 extending over a relatively narrow zone. Figures 4 (“blank result”) and 7 suggest that low probability values that spread out evenly through out the 3D block indicate of a “no leak” condition. As the released volume increases, probability values will increase in narrow zones and drop to near 0 elsewhere, thereby indicating zones of resistivity change associated with leaks.

Similarly, the voxel-wise average resistivity ratio images can be also used to infer the tank leakage (middle column of images in Figure 7). The right column of images shows an iso-surface view of the same results; resistivity ratios from 0.9 to 1.0 are transparently rendered. We suggest that this aspect of the posterior distribution allows an analyst to estimate the location, volume and shape of the contaminated zone. As expected, the voxel-wise average ratio decreases (i.e., the soil becomes more electrically conducting) from about 0.95 (top middle row, released volume = from 0.34 m^3) to about 0.2 (bottom middle row, released volume = 2.16 m^3). Also, the vertical and horizontal extent of the anomaly grows with increasing released volume. The sequence shown suggests that the plume represented by these inversions is behaving as expected (strong changes along a relatively narrow zone near the release point), thereby increasing our confidence in the reliability of the stochastic inversions.

Once the presence of a leak has been detected, tank managers may wish to plan mitigating or clean-up activities. Posterior distribution samples can also be used to estimate the probability of various parameters of interest. For example, information such as the volume of the anomaly or the location of the anomaly’s center of resistivity change may be useful (CRC—is computed using the equation for center of mass but substituting resistivity change for mass). Figure 8 shows the parameters’ estimated posterior distributions based upon the generated posterior image samples.. The left column plots show how the CRC moves with increasing released liquid volume. Early on during the spill sequence (released volume = 0.34 m^3), the most likely x and y coordinate position is about 14 m near the tank’s center (the release point is located at the tank’s center, $x = y = 14 \text{ m}$). As the released volume increases, the most likely x coordinate value changes to about 11.5m, indicating that the CRC is moving to the East. Likewise, the most likely y coordinate value changes to about 16 m, indicating that the CRC is moving slightly to the

South. The Z coordinate values moves from about 6 to about 8 m (the depth increases by 2 m).

Similar observations can be made regarding the volume of the invaded zone (VIZ – the volume of soil invaded by the plume: sum of all voxel volumes showing resistivity ratios indicating a change, i.e., values < 1.0). Early on during the release sequence (released volume = 0.34 m^3), the most likely VIZ value is about 120 m^3 . As the released volume increases, the likely VIZ value also increases, as expected. The 2.16 m^3 plot suggests that the most likely value is about 400 m^3 but there is a significant probability that it could be as low as 300 m^3 .

Convergence:

The output from our MCMC inversion approach provides a sample from the posterior distribution, which is the conditional probability distribution of the state of nature, given the data. The inversion results described so far are calculated only using samples generated after stationarity of the Markov process is achieved. All samples generated before that point are discarded. Here we discuss the techniques used to demonstrate that stationarity has been attained; i.e., the point in the simulation in which the samples generated are representative of the posterior distribution because they are unaffected by the starting point of the Markov chain(s) and have explored sufficiently the distributional structure.

In our application, the state of nature refers to a complicated, multi-dimensional configuration like the subsurface zone that is changing electrical properties over time. It may also refer to a particular related parameter of interest such the center of mass or the volume of the region. The posterior distribution can be thought of as the best stochastic description of the state of nature that incorporates all pertinent physical and theoretical models as well as observed data. Characterization of the posterior distribution is the primary goal in the Bayesian statistical paradigm. In our applications, however, analytical calculation of the posterior distribution is impossible, and it is only feasible to draw a representative sample from the distribution of interest. Asymptotic results ensure that the sample ultimately spans the entire posterior distribution and reveals the actual state frequencies that characterize the posterior. In statistical terms, the sample forms an ergodic Markov chain with a stationary distribution $P(m)$. This means that once the chain has taken a sufficient number of steps, T_0 , the (unconditional) distribution of the state, $x^{(T)}$, at any step $T \geq T_0$ (i.e., is “stationary”, collecting more samples will not change the posterior distribution), and is $P(m)$. We call T_0 the “burn-in” period. Once burn-in has occurred, a chain has essentially “forgotten” where it started, stationarity has been achieved, and the remaining samples are from the desired posterior distribution.

The diagram on the left side of Figure 9 schematically illustrates some of the key ideas pertaining to convergence. The area within the rectangle represents the space of possible solutions, S . The three curves represent the route followed by three Markov chains, each one having a different starting point. The posterior distribution region is located within the dashed elliptical region. The bull’s eye symbols identify the location of two modes in the posterior distribution. All samples collected within the burn-in period are discarded

because the chain's route is still influenced by the chain's starting point. After burn-in, the Markov chains have converged when the route followed by all three chains is unaffected by the chains' starting points, all have sampled both modes, and indicate similar heights for the mode peaks.

The plots on the right side of Figure 9 illustrate two types of metrics that: a) identify the "burn-in" period and determine the Markov chain segment that is stationary, and b) confirm stationarity by a normality assessment. Glaser (2003) provides details pertaining to these metrics; here we summarize some of the key points. The top plot (Gelman-Rubin diagnostic) is a heuristic metric that uses multiple parallel Markov chains to simultaneously estimate the burn-in period length T_0 and corroborate the claim that remaining samples come from a stationary distribution. It is based on the idea that different Markov chains having different starting points share a common, but unknown, limiting distribution, the posterior $P(m)$. When the metric levels off to a value near 1.0 and approximately remains at this level, it suggests that the burn-in period ended at iteration T_0 ; posterior samples are being accumulated from this point on. Also, stable values near 1.0 suggest that samples after T_0 are an adequate characterization of the posterior distribution, as exploration of the parameter space has apparently succeeded in visiting all the significant distributional modes. For the example shown, the burn-in period appears to end around iteration $T_0 = 800$.

The bottom plot in Figure 9 provides confirmatory evidence that the posterior is being sampled by showing that averaged samples after burn-in are normally distributed. When the central limit theorem is applied to MCMC algorithms, it asserts that the value of a parameter averaged over a large number of iterations after burn-in is approximately normally distributed. Thus, when samples from a stationary Markov chains are used, samples of a given parameter of interest (in our case, the centroid of the changing region) should be normally distributed. Larger p -values provide stronger evidence of normality. A p -value below 0.05 is generally considered reason to question the validity of the assertion of normality. The lower plot in Figure 9 provides strong evidence of normality. Taken together, the plots suggest that the Markov chains in this example have converged to the limiting posterior distribution $p(m)$, and that we can reliably employ these posterior samples to estimate the parameter(s) of interest.

Posterior Analysis -- Clustering

The inversion approach described here produces a large number of samples from the posterior distribution. In most applications, it is necessary to summarize the relevant information in the posterior distribution so that it can be visualized and understood. In a state space M , there are hills and valleys in the likelihood surface, their height is proportional to the probabilities for each of the states \underline{m}_j , and the location of a given peak is commonly referred to as a mode. The posterior analysis method used to produce previous images (voxel-wise averages) provides useful summary of samples from a uni-modal posterior distribution. But, when the posterior distribution is multi-modal, this approach can produce misleading results.

We now consider an approach based on clustering instead of the simple voxel-wise averages described earlier. In the 1-D case, a simple histogram plot reveals if the data have come from a uni-modal distribution or not. In the 2-D case also, a surface plot will reveal the same information although a clustering operation will reveal the location and degree of concentration about the modes in both. Images are highly multidimensional because each voxel corresponds to a dimension in the sample space resulting in a space of N dimensions where the number of voxels D may be typically in the thousands, if not in the tens of thousands (e.g., an image with $25 \times 25 \times 25$ voxels has 15625 dimensions).

We extend the same idea to dimensions much larger than two, in order to understand the nature of the posterior distribution directly from the sample images. When a distribution is determined to be uni-modal based on the cluster analysis one can try to estimate the different moment parameters such as mean, co-variances etc. For a multi-modal distribution on the other hand these parameters even though estimable, will lose their intuitive contents. Nevertheless, it is still possible to empirically analyze each cluster as a uni-modal distribution. For instance, in our example we may look at the frequency distribution of the within-cluster distances defined as the distances from the cluster members to the cluster center.

The problem of posterior inference is further complicated when we have discrete valued image samples as is the case for the resistivity image space. One way to manage the inference problem in these high-dimensional cases is to compute parameters of topical interest in the image. For example, one may consider a feature space generated from an image sample where the first set of features may consist of the number and sizes of the connected regions of each of the possible resistivity values. The next set may give the (x, y) coordinates of the location of their centroids and so on. Depending on the questions of interest, these may give adequate information related to the samples. In other cases however, the experimenter may be interested in the modes of the posterior distribution of the sample space itself, where the samples are multidimensional N vectors and then treating the mode-modal frequency pairs as a summary description of the posterior distribution of the features under study.

We now discuss another methodology of posterior analysis that is suitable for multi-modal and multi-dimensional *category-valued* distributions. In this approach, clusters along with their sample space locations and frequencies are interpreted as meta-states. A detailed description of this approach is presented in Sengupta and Ramirez (2003). Here we summarize some of the key points. Clustering is a methodology focused upon grouping samples from a distribution based on the similarity of the samples themselves, or on the similarities of carefully chosen sample features. It is appropriate for the analysis of multi-dimensional spaces and multi-modal distributions. Our clustering algorithm attempts to group together samples that are similar to one another, thereby providing a good summary of the posterior distribution under study. The clustering process is implemented by measuring the distance d in state space between two given images or between an image and a cluster center. The former simply expresses how much global dissimilarity exists between two sample images whereas the distance d in the later case is defined in terms of the distance of the sample from the ‘cluster histograms.’ A cluster

histogram consists of an array of distributions with each voxel distribution constituting an element of the array. Each voxel distribution in turn is based on the frequency distribution of different categories within the cluster samples for that voxel. In terms of the image samples with categorical value, this entails the computation of a global measure of d between an image and a cluster of images that tells one how much overall dissimilarity exists between the two. In computing the sample-cluster ‘distance’, we look at the global (including all pixels or volume elements) marginal distribution vector and compute how similar a given image sample is compared to the distribution of the voxels in a cluster.

This algorithm will find the cluster zones and their corresponding frequency(ies) of the (possibly) multi-modal distribution. The cluster frequency is proportional to the cluster probability because the Metropolis Hastings algorithm samples the state space at a rate that is proportional to the posterior probabilities. We use the cluster location (i.e., the mean in the continuous case or the mode in the discrete case) along with the cluster frequency as the best summary representation for each cluster. In our approach the number of clusters is allowed to vary during clustering based on the internal configuration of elements in the clusters as well as the distance between adjacent clusters in state space. This approach allows the experimenter the opportunity to examine the most likely configuration of the resistivity structure at a given location, ranked according to their probabilities.

We used this clustering approach to analyze some of the posterior distributions described earlier. Figure 10 shows the cluster analysis results for the sand-lead target experiment. The upper left image shows a vertical slice through the target. The upper right image shows the mean resistivity ratio obtained when a uni-modal posterior distribution is assumed. The bottom row of images shows the results of clustering analysis assuming a multi-modal distribution. These are the cluster centroids (mean resistivity ratio) for the three most likely (three highest frequencies) clusters. The frequency values have been normalized relative to the number of posterior distribution samples (4302) to obtain relative frequency values for each cluster. The bottom left image represents the centroid for the most frequent cluster: four out of every five samples come from this part of the state space. Note that the location, and size of the cluster centroid are similar to the target. The maximum value resistivity ratio magnitude is about 2.48 while the target value is 2.51. The middle and right bottom row images represent centroids for low frequency clusters. The middle image shows two closely spaced anomalies that straddle the target’s elevation and a resistivity ratio value (16.0) that is substantially larger than the target. The bottom right image also shows an anomaly shape that is significantly different from the target and has a maximum resistivity ratio value (16.0) that is substantially larger than the target.

We believe that the results shown in Figure 10 illustrate the value of multi-cluster analysis. When we perform a multi-modal analysis, the resistivity ratio for the most frequent cluster is about 2.48, close to the target value of 2.51. If we assume a uni-modal distribution, the value is about 3.0, a poorer match to the target, because all samples have lumped into a single cluster. The multi-cluster analysis segregates posterior samples with

lower and higher values into separate clusters, thereby improving the accuracy of the resistivity value corresponding to the most frequent cluster (bottom left cluster in Figure 10). Similar comments apply to the size, location and shape of the cluster anomalies.

Cluster analysis allows us to analyze the variability and uncertainty that is supported by the available data. It provides the basis for the objective analysis of alternative models when the available information isn't sufficient to uniquely identify the system under study. Furthermore, the consequences due to sources of uncertainty such as measurement error, and lack of sensitivity or resolution are explicitly analyzed with by this approach.

The tank leak results were also analyzed using our clustering approach. Figure 11 shows the results obtained for the case where 2160 liters had been released. The left column of graphs displays the three most likely cluster centroids with frequencies ranging from 0.12 to 0.38. These three clusters encompass 68 % of all posterior distribution samples. We note that all of them suggest a roughly vertical anomaly directly below the release point to depths ranging from 8 – 10 m.

The clusters in Figure 11 helps the analysts identify alternate models that are consistent with available data and their corresponding probabilities. The left column of frames show the voxel-wise mean resistivity ratio for the top three most probable clusters. The frequency values shown have been normalized relative to the number of posterior distribution samples (3802) to obtain relative frequency values for each cluster. Clusters A and B suggest that a strong vertical anomaly exists just East of the release point. Cluster A suggests the possibility of liquid ponding to the West of the release point at a depth of about 8m. Likewise, cluster C suggests a wider but somewhat weaker anomaly. An analyst could consider cluster A as the most likely thereby inferring a pillar-shaped vertical invasion zone with possible liquid ponding to the West of the release point at 8m depth. Considering clusters A and B, the analyst may also infer that the bulk of the contamination lies below and to the East of the release point.

The right column of frames in Figure 11 shows the center states for the clusters shown. These center states are those posterior samples having the smallest distance d to the cluster's center point. We can use the centers states as a way to evaluate a cluster's central tendencies. These central tendencies should be unaffected by the variability of such properties within a cluster; clearly, this variability affects the mean values shown in the left column of frames in Figure 11. The center states for clusters A and B show similar location and shape as the cluster mean image. However, the mean for cluster A suggests the possibility of liquid ponding to the West of the release point at a depth of about 8m whereas the cluster center state does not show such features. This means that the possible ponding of liquid is not part of the central tendency for cluster A, but is a part of some of the states around the cluster center. The center state for cluster B suggests a stronger resistivity change (0.03) than the values near 0.2 indicated by the cluster's mean.

The clustering analysis described here allows an analyst to consider alternate hypotheses that are consistent with all available data, and their corresponding probabilities. Tank

operators can take into account these alternate hypotheses when deciding among environmental management alternatives that mitigate contaminant spreading or clean-up the subsurface.

Computational expense – parallel computing

The MCMC method we have used is computationally expensive. For example, a tank leak problem involving 28800 voxels and 128 transmitting electrodes and a typical workstation with one CPU, will require about 45 days to accumulate a sufficient number of posterior samples (about 4000). Almost all of that time is used to solve the forward problem. In comparison, a deterministic inversion of the same problem using the same workstation may take 8 – 12 hours.

Clearly, parallel computation of the forward problem is required to make the MCMC approach practical. We have parallelized the problem in two ways: 1) each of the individual Markov chains is run on separate processors. 2) the computational load for each chain is further distributed amongst multiple processors by computing the potential field due to each transmitter electrode on separate processors. The potential fields are then post-processed to compute the transfer resistances associated with individual readings. When 128 processors are used, this approach reduces the processing time to about 12 hours. Processing times of 12 hours are acceptable for many real-life applications of the MCMC approach.

Conclusions:

We have discussed a stochastic methodology for Bayesian inversion of changing subsurface electrical resistivity data. This methodology, called the Stochastic Engine, is based upon Bayesian inference and is implemented via a Markov Chain Monte Carlo algorithm. The inversion of electrical resistivity data is an ill-posed problem requiring regularization. Our approach makes use of prior information to sufficiently reduce the size of the space of feasible solutions in order to mitigate ill-posedness. The base representation consists of up to ten blocky regions of resistivity change embedded within an unchanging volume. Additional information can include the sense of the change (increasing or decreasing resistivity), upper/lower bounds for the volume of the changing region, resistivity change magnitude and spatial relationships of the trapezoidal regions (e.g., requiring partially overlapping or contiguous blocks).

We have tested the performance of the SE algorithm using a series 3D physical models and a field experiment. The performance tests indicate that the algorithm locates the most likely location(s), shape(s), and volume of the changing region(s). This Bayesian reconstruction technique appears to produce robust results when noisy field data is used. The method also allows the use of disparate data types (e.g. measurements of the volume of the liquid infiltrating the subsurface, geophysical logs indicating changing conditions along a borehole) to constrain the posterior probabilities.

In contrast to conventional methods of inverting for resistivity change, the aim of the SE methodology is to estimate the a-posteriori distribution defined over the entire space of feasible solutions instead of simply finding a single estimate of the unknown resistivity

distribution. The SE approach is particularly useful when prior information is available, alternate models consistent with the data are desired, and measures of resolution and solution uncertainty are needed.

We have briefly discussed metrics that indicate when the samples generated from the MCMC simulation are truly representative of the posterior distribution and hence capable of supporting inference. . We have also discussed a categorical clustering technique that partitions posterior distribution samples into multiple clusters whose members have similar characteristics. Cluster analysis allows the identification of multiple modes (i.e., alternative system configurations that are consistent with available data) and their associated probabilities.

The method is computationally intensive. For example, the Hanford results previously discussed required computation times of approximately 12 hours using a 128 CPU parallel machine. The availability of large parallel computers greatly reduces the required processing times; making the approach more practical for real life applications.

We view the SE and the deterministic inversion methodologies as complementary approaches. The deterministic method is likely to be the preferred method when little or no prior data is available, a single data type is being inverted, and when fast inversion times are required. The SE method is most useful when inverting problems with many secondary extrema, contradictory or sparse data, with multiple data types, and problems having useful prior information. The SE method is also a good choice for problems where the linearization of a non-linear problem must be avoided, explicit estimates of solution uncertainty are required, and alternate models that consistent with the data are desired.

Acknowledgments

This work was funded by the Laboratory Directed Research and Development Program at Lawrence Livermore National Laboratory. DOE's Office of River Protection, CH2M Hill Hanford Group sponsored the leaky tank experiment. The experiment was designed and managed by G. Gee, B. Barnett, and M. Sweeney, Pacific Northwest National Laboratory. The physical model data was collected by Stan Martins (LLNL) as part of a Laboratory Directed Research and Development project led by Charles Carrigan (LLNL). We would like to thank Dr. Andrew Binley, Lancaster University (UK) who served as an informal reviewer of this manuscript. This work was performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract W-7405-ENG-48.

References:

Aines, R., J. Nitao, R. Newmark, S. Carle, A. Ramirez, D. Harris, J. Johnson, V. Johnson, D. Ermack, G. Sugiyama, W. Hanley, S. Sengupta, W. Daily, R. Glaser, K.

- Dyer, Y. Zhang, Z. Yu, R. Levine, 2002,. *The Stochastic Engine Initiative: Improving Prediction of Behavior in Geologic Environments We Cannot Directly Observe*, UCRL-ID-148221, Lawrence Livermore National Laboratory, Livermore, CA
- Barnett, B., G. Gee, and M. Sweeney, 2002,. Results of tank leak detection demonstration using geophysical techniques at the Mock Tank test site, PNNL-13818, Pacific Northwest National Laboratory, Richland, WA
- Binley, A., S. Henry-Poulter and B. Shaw, 1996 Examination of solute transport in an undisturbed soil column using electrical resistance tomography, *Water Resour. Res.*, **32**(4), 763-769.
- Binley, A., P. Winship, M. Pokar and J. West, 2001, Cross-borehole radar and resistivity tomography: A comparison of techniques in unsaturated sandstone, In: Proc. Symp. Applications of Geophysics to Engineering and Environmental Problems (SAGEEP2001), Environmental and Engineering Geophysical Society, Denver, CO.
- Andersen, K.E., Brooks, S.P. and Hansen, M.B. 2001a, Bayesian inversion of geoelectrical resistivity data. Research Report R-01-2016. Department of Mathematical Sciences, Aalborg University, Denmark.
- Andersen, K., S. Brooks, and M. Hansen, 2001b, A Bayesian approach to crack detection in electrically conducting media, *Inverse Problems*, vol. 17, p121-136.
- Daily, W. A. Ramirez, D. LaBrecque and J. Nitao, 1992, Electrical resistivity tomography of vadose water movement, *Water Resour. Res.*, **28**(5), 1429-1442.
- Daily, W. D., W. Lin and T. Buscheck, 1987, Hydrological Properties of Topopah Spring Tuff: Laboratory Measurements, *J. Geophys. Res.*, **92**, 7854-7864.
- Daily, W., A. Ramirez, A. Binley, D. LaBrecque, 2001, Electrical Resistance Tomography – Theory and Practice, in *Near Surface Geophysics*, Volume 2, Society of Exploration Geophysics (in press); also in UCRL-JC-144936, Lawrence Livermore National Laboratory, Livermore, Ca.
- Ellis, R.G. and D.W. Oldenburg, 1994, The pole-pole 3-D DC-resistivity inverse problem: a conjugate gradient approach, *Geophys. J. Int.*, **119**, 111-119.
- Fox, C. and Nicholls, G. 'Sampling Conductivity Images via MCMC'. In: The Art and Science of Bayesian Image Analysis, K.V. Mardia, C.A. Gill, R.G. Aykroyd eds, Proceedings of the Leeds Annual Statistical Research Workshop (LASR), 91-100, July 1997, Leeds University Press, 1997.

- Gee, G. and, A. Ward, 2001, Vadose Zone Transport Field Study: Status Report, Pacific Northwest National Laboratory, PNNL-13679, Richland, WA.
- Glaser, R., C. Lee, J. Nitao, and W. Hanley, 2002, A Markov Chain Monte Carlo based method for system identification, UCRL-JC- **150494**, Lawrence Livermore National Laboratory, Livermore, CA.
- Glaser, R., 2003, Stochastic engine convergence diagnostics, UCID-??? Lawrence Livermore National Laboratory, Livermore, CA.
- Hearst, J., P. Nelson, and F. Paillet, 2000, *Well logging for physical properties*, John Wiley & Sons, New York.
- Kaipio, J., V. Kolehmainen, E. Somersalo, and M. Vauhkonen, 2000, Statistical inversion and Monte Carlo sampling in electrical impedance tomography, *Inverse Problems*, vol. 16, p 1487-1522.
- Kemna, A., A. Binley, A. Ramirez and W. Daily, 2000, Complex resistivity tomography for environmental applications, *Chemical Engineering Journal*, **77**(1-2), p11-18.
- Kolehmainen, V., E. Somersalo, P. Vauhkonen, M. Vauhkonen, and J. Kaipio, 1998, A Bayesian Approach and total variation priors in 3D electrical Resistance Tomography, In: *Biomedical Engineering Towards the Year 2000 and Beyond. Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Hong Kong, China, October 29 - November 1, 1998, p. 1028-1031.
- LaBrecque, D. J., Miletto, M., Daily, W. D., Ramirez, A. L., and Owen, E., 1996, The effects of noise on Occam's inversion of resistivity tomography data, *Geophysics*, **61**, 538-548.
- LaBrecque, D., G. Morelli, W. Daily, A. Ramirez, and P. Lundegard, 1999, Occam's inversion of 3D Electrical Resistivity Tomography, in "Three-dimensional electromagnetics", eds. M. Oristaglio, B. Spies, and M.R. Cooper, Soc. Of Exploration Geophysicists, pages 575-590.
- LaBrecque, D.J. and Yang, X., 2000, Difference inversion of ERT data: A fast inversion method for 3-D in-situ monitoring, In: *Proceedings of the Symposium on the Application of Geophysics to Engineering and Environmental Problems*, Env. Eng. Geophys. Soc., 907-914.
- Metropolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller, 1953, Equation of state calculations by fast computing machines, *J. Chem. Phys.*, **1**, no. 6, p. 1087-1092.

- Morelli, G., and LaBrecque, D. J., 1996, Advances in ERT modeling, *Eur. J. Environ. Eng. Geophys.*, **1**, 171-186.
- Mosegaard, K., and A. Tarantola, 1995, Monte Carlo sampling of solutions to inverse problems, *Journal of Geophysical Research*, vol. 100, no. B7, p12431-12447.
- Mosegaard, K., 1998, Resolution analysis of general inverse problems through inverse Monte Carlo sampling. *Inverse Problems*, v. 14, pp. 405-426.
- Mosegaard, K., and M. Sambridge, 2002, Monte Carlo analysis of inverse problems. *Inverse Problems*, v. 18, pp. R29 - R54.
- Mosegaard, K., and C. Rygaard-Hjalsted, 1999, Probabilistic analysis of implicit-inverse problems. *Inverse Problems*, v. 15, pp 573-583
- Newmark, R., R. Aines, J. Nitao, W. Hanley, S. Carle, A. Ramirez, S. Sengupta, and D. B. Harris, 2001. *Stochastic Engine: Direct Incorporation of Measurements Into Predictive Simulations*. Proceedings of the International Groundwater Symposium "Bridging the gap between measurement and modeling in heterogeneous media" Berkeley, CA March 25-28, 2002.
- Park, S.K. and G.P. Van, 1991, Inversion of pole-pole data for 3-D resistivity structure beneath arrays of electrodes, *Geophysics*, **56**, 951-960.
- Ramirez, A., W. Daily, D. LaBrecque, E. Owen and D. Chesnut, 1993, Monitoring and underground steam injection process using electrical resistance tomography, *Water Resources Research*, **29(1)**, 73-87
- Ramirez, A., W. Daily, A. Binley, D. LaBrecque and D. Roelant, 1996, Detection of Leaks in Underground Storage Tanks Using Electrical Resistance Methods, (UCRL-JC-122180, October, 1995), *J. Engineering and Environmental Geophysics*, **1**, 189-203.
- Ramirez, A., and W. Daily, 2001, Electrical Imaging of the Large Block Test-Yucca Mountain, Nevada, *Journal of Applied. Geophysics* , **46**, 85-100.
- Rodgers, C., 1976, Retrieval of atmospheric temperature and composition from remote measurements of thermal radiation, *Reviews of Geophysics and Space Physics* , 14, no. 4, 609-624.
- Rodgers, C., 1990, Characterization and error analysis of profiles retrieved from remote sounding measurements, *Journal of Geophysical Research*, 95, no. D5, 5587-5595.
- Sasaki, Y, 1994, 3-D resistivity inversion using the finite-element method, *Geophysics*, **59(11)**, 1839-1848.

- Sengupta, S., and A. Ramirez, 2003, Exploratory global inference from posterior image samples in a Markov chain Monte Carlo simulation experiment: A high-dimensional data clustering approach with extensions to categorical data, Lawrence Livermore National Laboratory, UCRL-JC-????, Livermore CA
- Shima, H., 1992, , 2-D and 3-D resistivity image reconstruction using crosshole data, *Geophysics*, **57**, 1270-1281.
- Slater, L., A. Binley, W. Daily and R. Johnson, 2000, Cross-hole electrical imaging of a controlled saline tracer injection, *J. Applied Geophysics*, **44**, 85-102.
- Ward, A. and, G. Gee, 2001, Vadose Zone Transport Field Study: FY 2001 Test Plan, Pacific Northwest National Laboratory, PNNL-13451, Rev. 1, Richland, WA.
- Yang, X, D. LaBrecque, 1999, Comparison between Stochastic and Occam's Inversion of 3-D ERT data. Proceedings of the Symposium on the Application of Geophysics to Engineering and Environmental Problems, Oakland, CA, March 14 – 18, 1999.
- Zhang, J., Mackie, R.L. and Madden, T., 1995, 3-D resistivity forward modeling and inversion using conjugate gradients, *Geophysics*, **60**(5), 1313-1325.

Figures

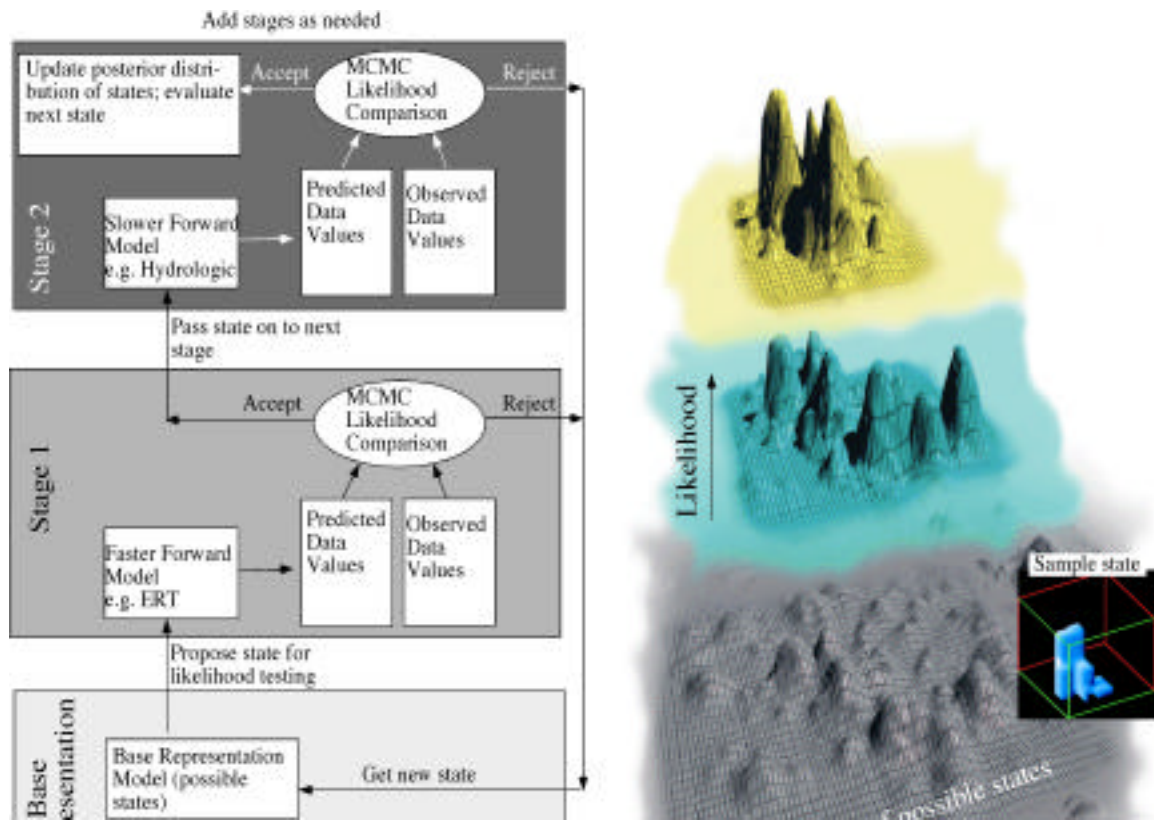


Figure 1. The left side shows a schematic diagram of our stochastic inversion algorithm. The method uses multiple data types, arranged in stages. The right side shows the impact that the MCMC search has on the posterior distribution's likelihood surface.




Figure 2 shows a schematic representation of the physical model set up. Four vertical electrode arrays were immersed in a fiberglass water tank. Various objects were inserted between the arrays at a variety of locations and ERT data collected.

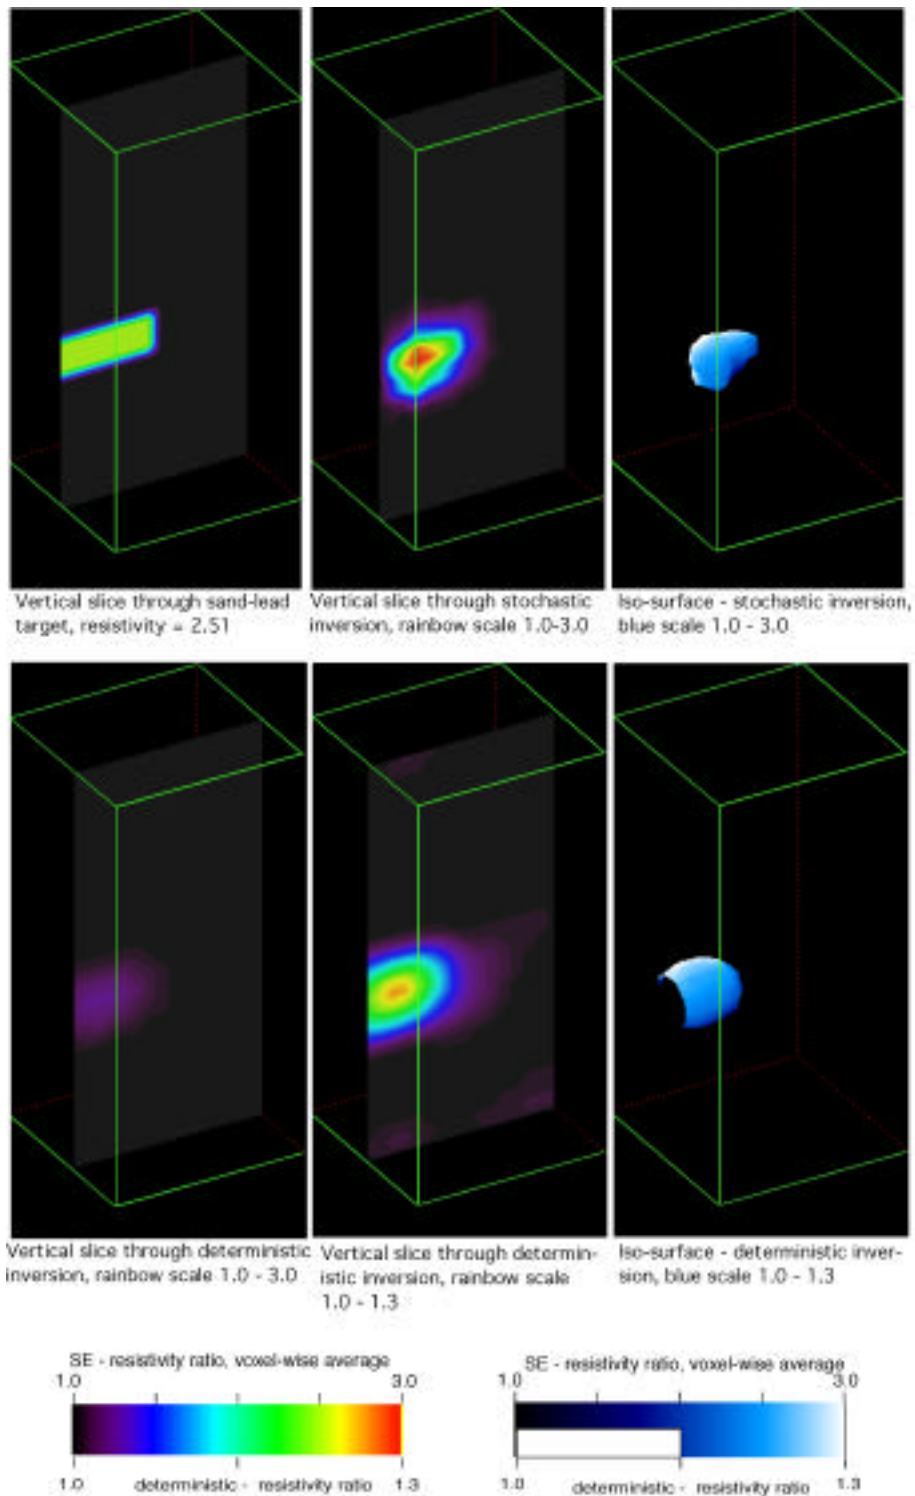


Figure 3 shows a comparison of stochastic and deterministic inversions. The target configuration is shown in the upper left frame, and the stochastic result shown by the top middle and right frames. The bottom row of frames shows the deterministic result. The white bar overlaying the blue color bar indicates the range of transparent values used to construct the iso-surfaces.

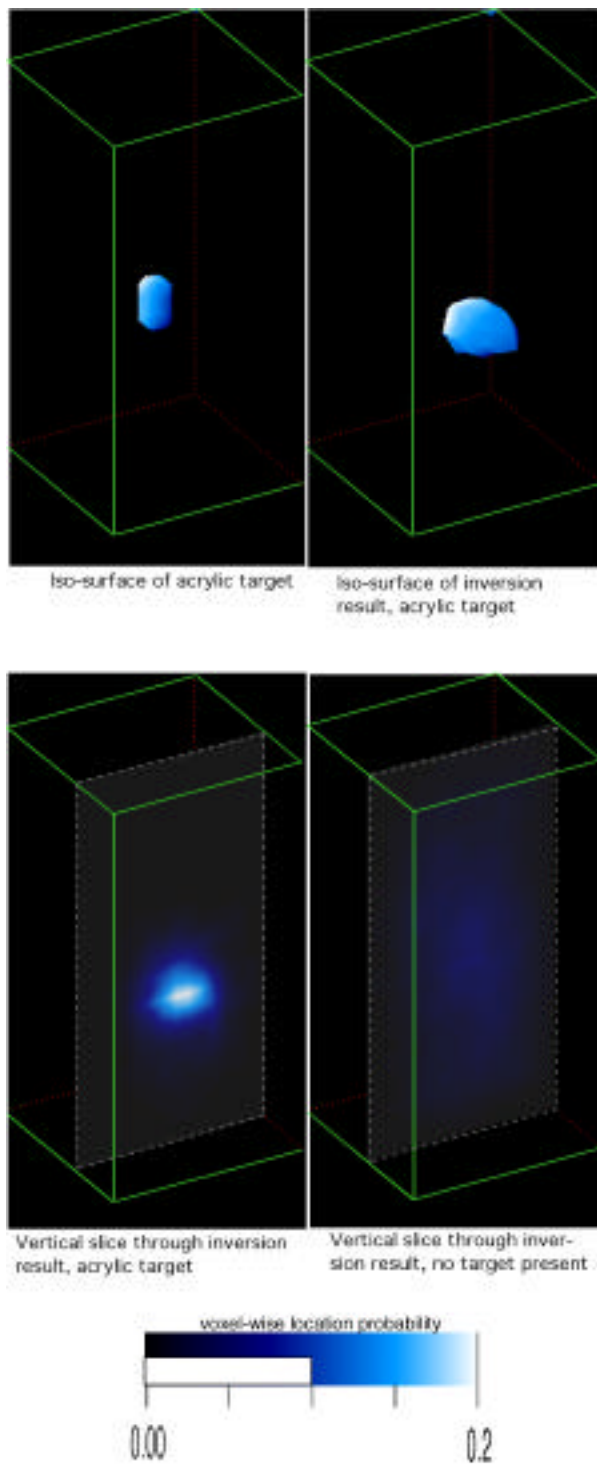


Figure 4 shows stochastic inversion results for the cases where the acrylic target is present and absent. The top two frames show iso-surfaces of the target and the corresponding inversion result. The bottom frames show vertical slices through the inversion block for cases of the target present (bottom left) and absent (bottom right). The white bar overlaying the blue color bar indicates the range of transparent values used to construct the iso-surfaces.




Figure 5 shows a schematic layout of the leaking tank site. Hypersaline brine solution was released from a point near the center of the tank's bottom. Sixteen vertical arrays of electrodes were used to monitor the infiltration process.

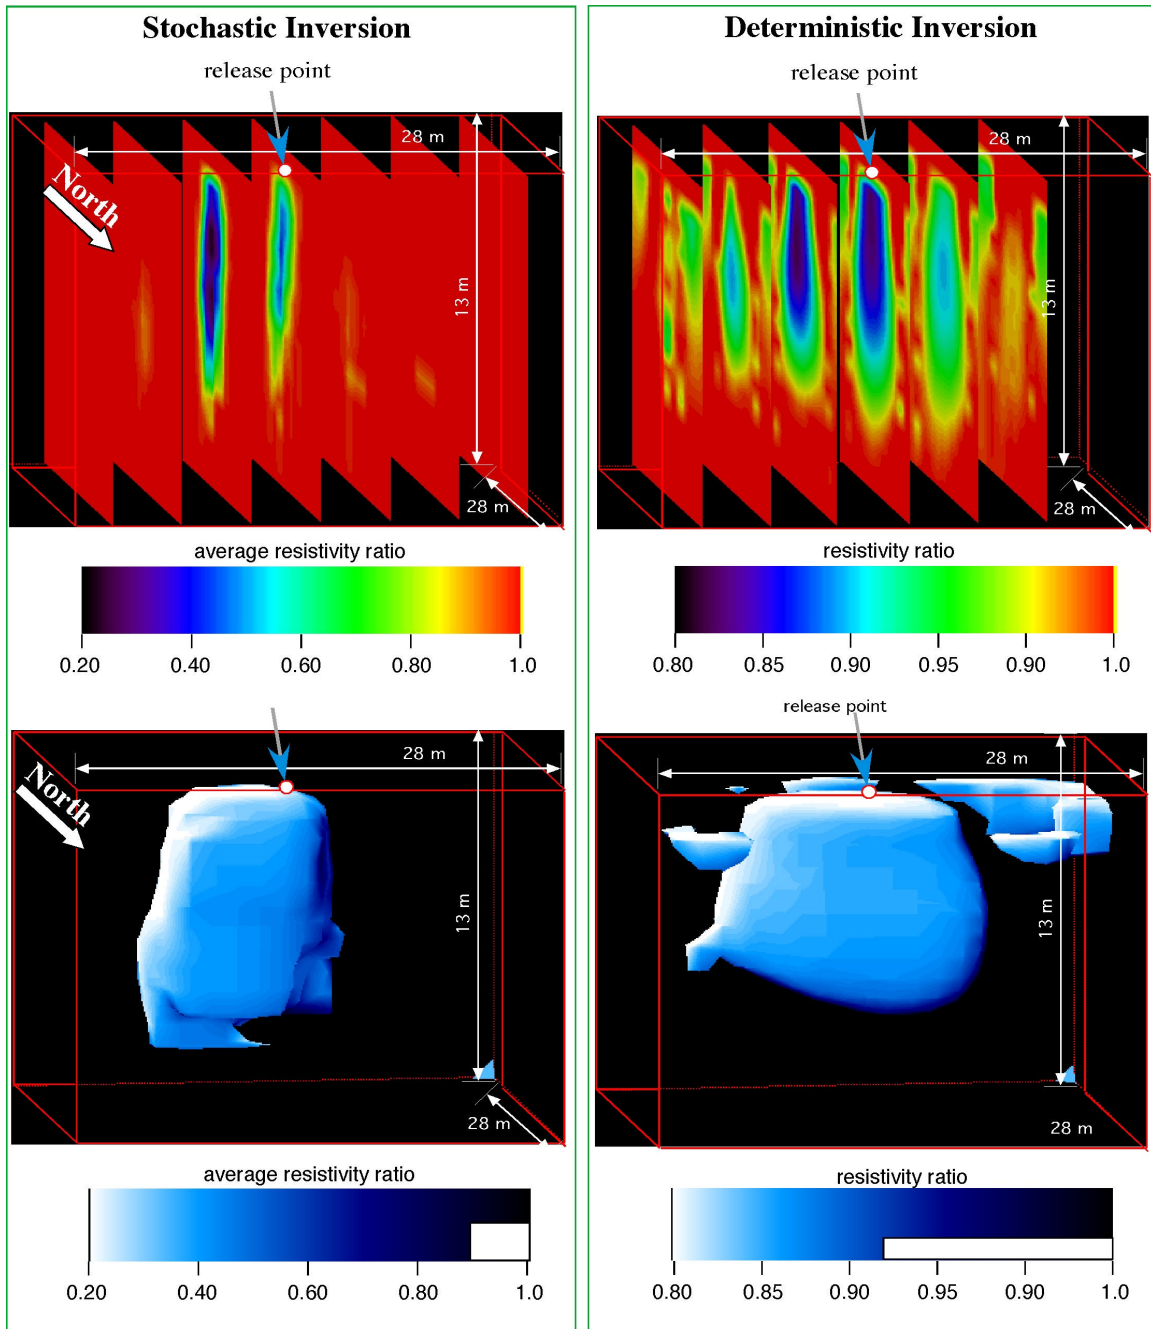


Figure 6 compares stochastic (left frames) and deterministic inversions (right frames). The top row of images consists of a series of vertical slices oriented parallel to North-South line. The bottom row shows iso-surfaces). The white bar overlaying the blue color bar indicates the range of transparent values used to construct the iso-surfaces.

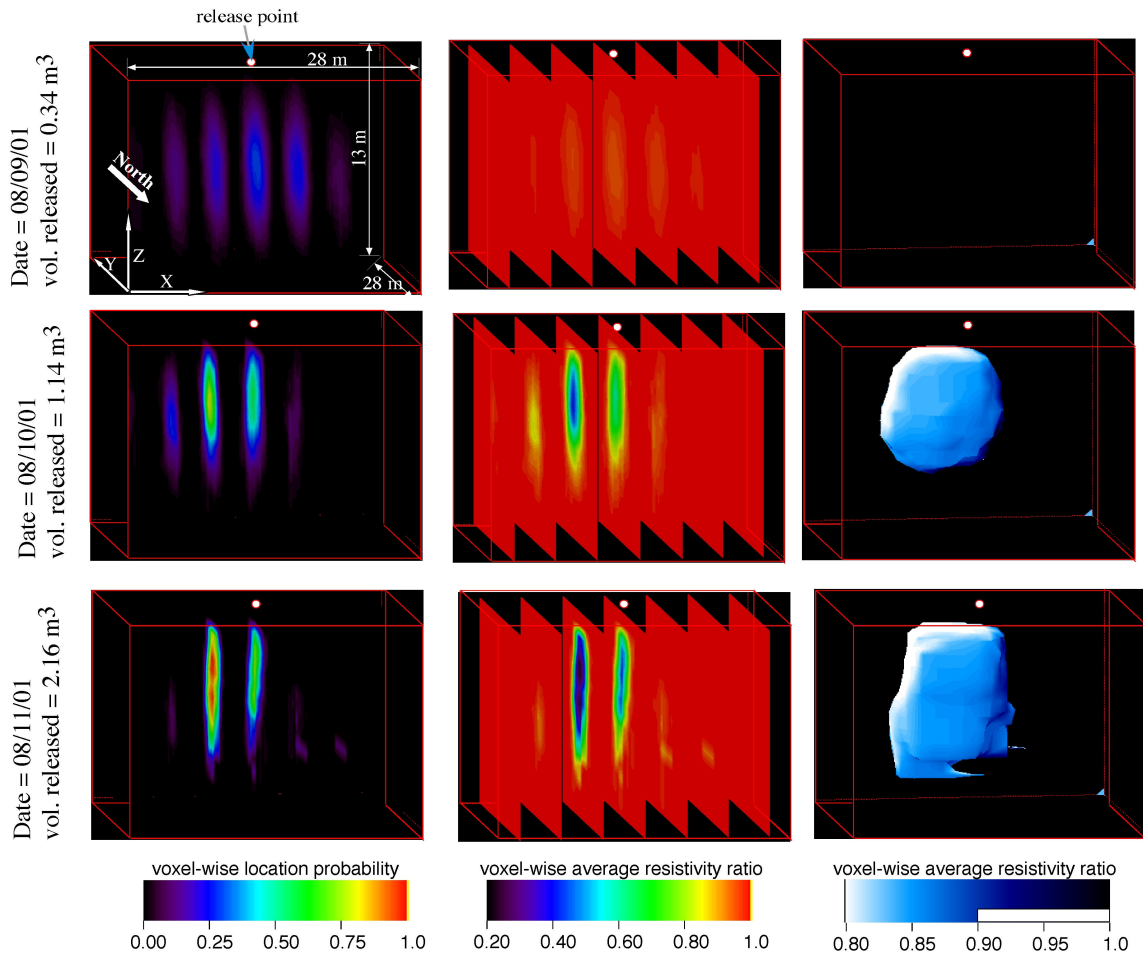


Figure 7 displays stochastic results obtained as a function of brine volume released. The left column shows voxel-wise location probabilities. The middle column shows the voxel-wise average resistivity ratio, and the right column of images shows the same results as iso-surfaces.

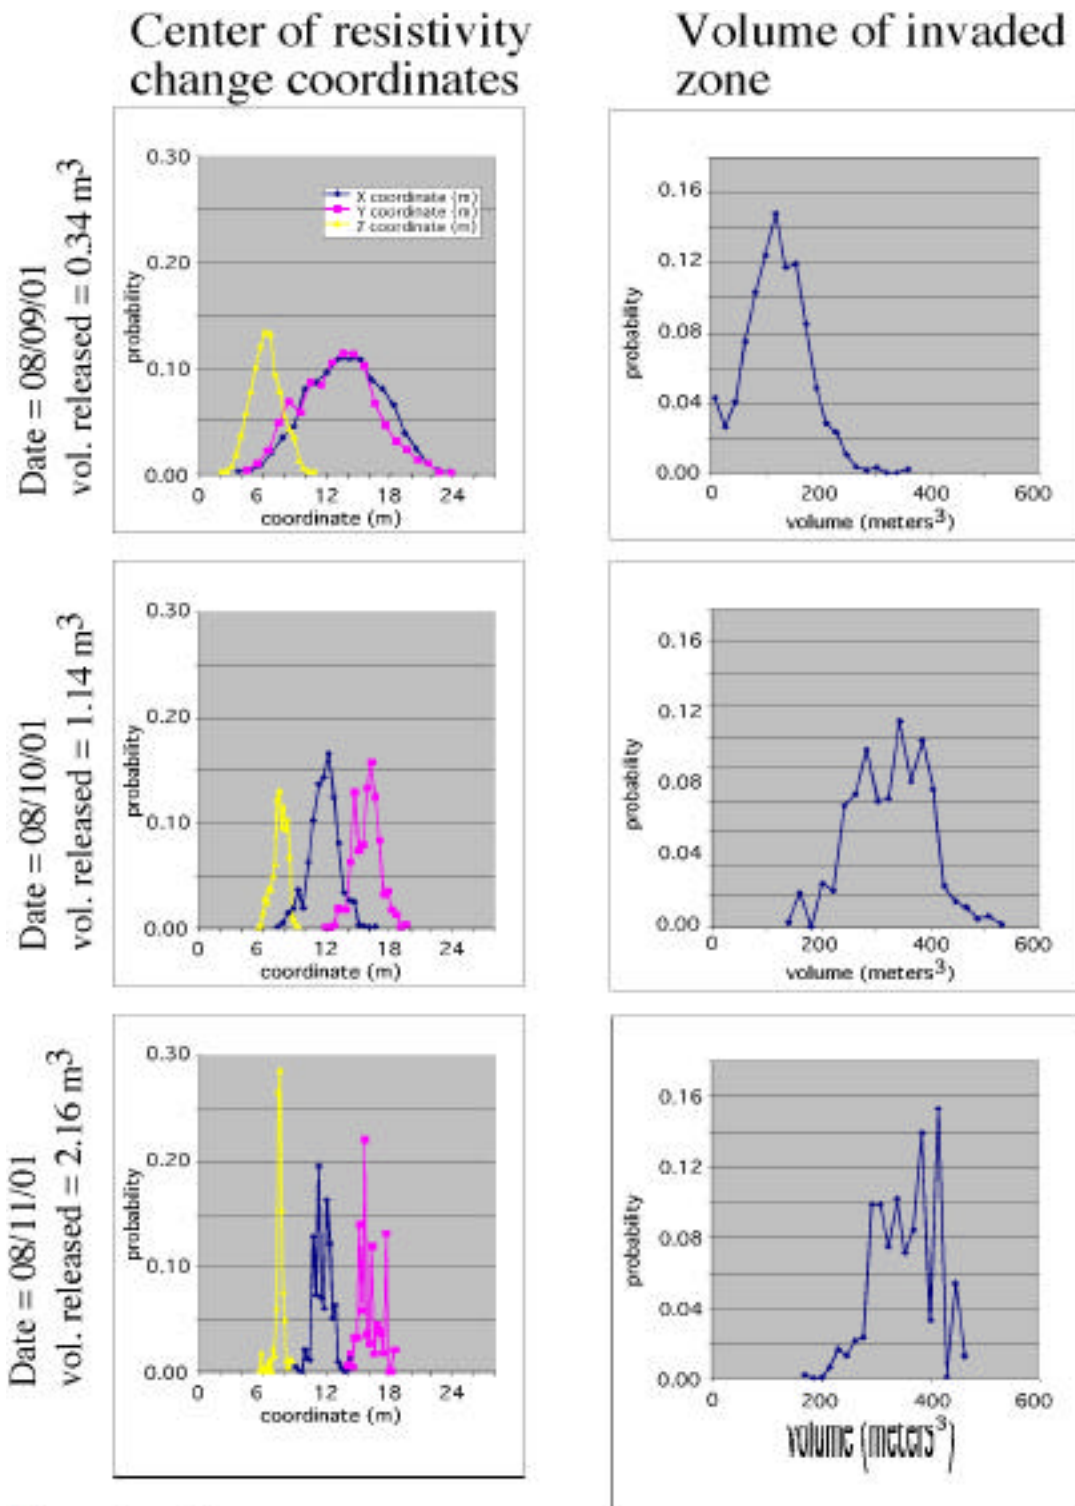


Figure 8 presents histograms of two parameters of interest as function of released brine volume. The left column of graphs shows the probability distribution for the coordinates of the center of resistivity change (equivalent to a center of mass). The right column shows the probability distribution of the volume of the invaded zone.

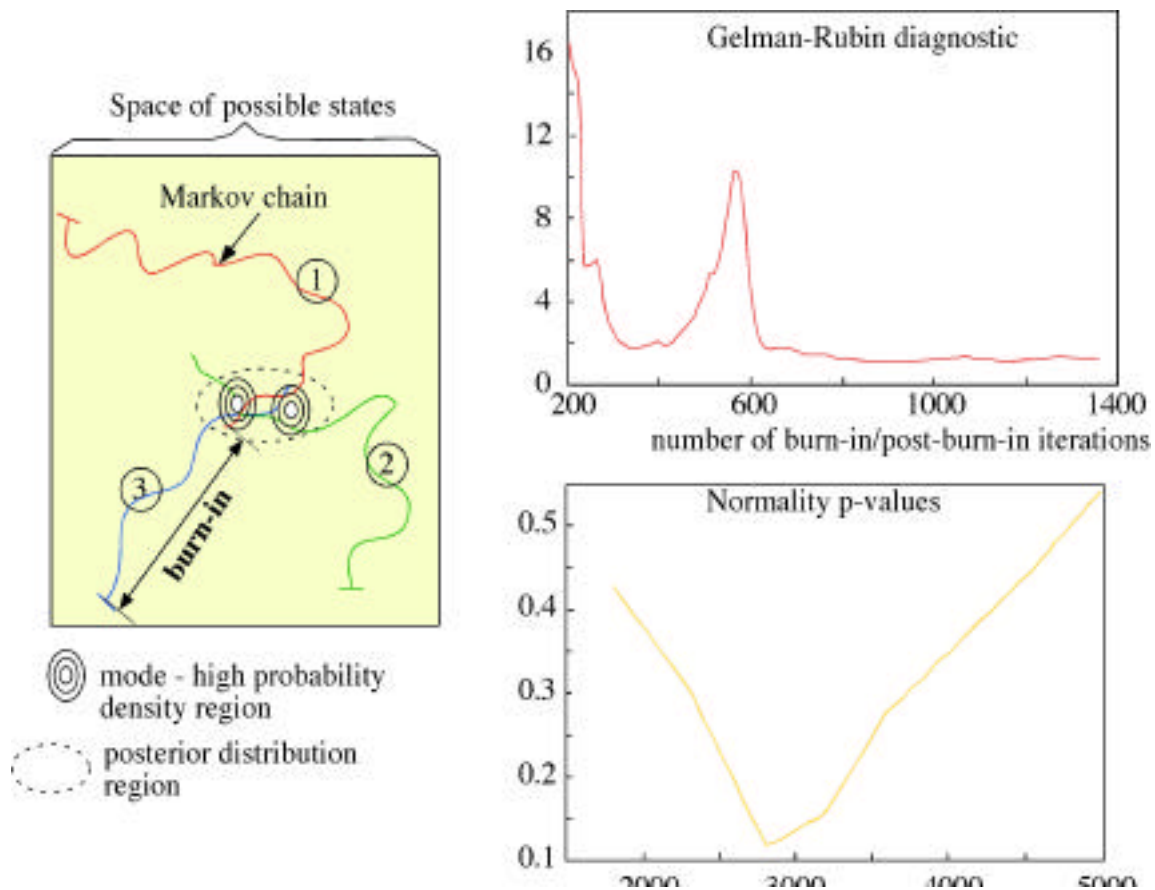


Figure 9. The right side shows a schematic diagram illustrating some of key ideas pertaining to Markov chain convergence: burn-in period and all chains sampling all posterior distribution modes. The left side shows plots of two metrics used to assess Markov chain convergence.

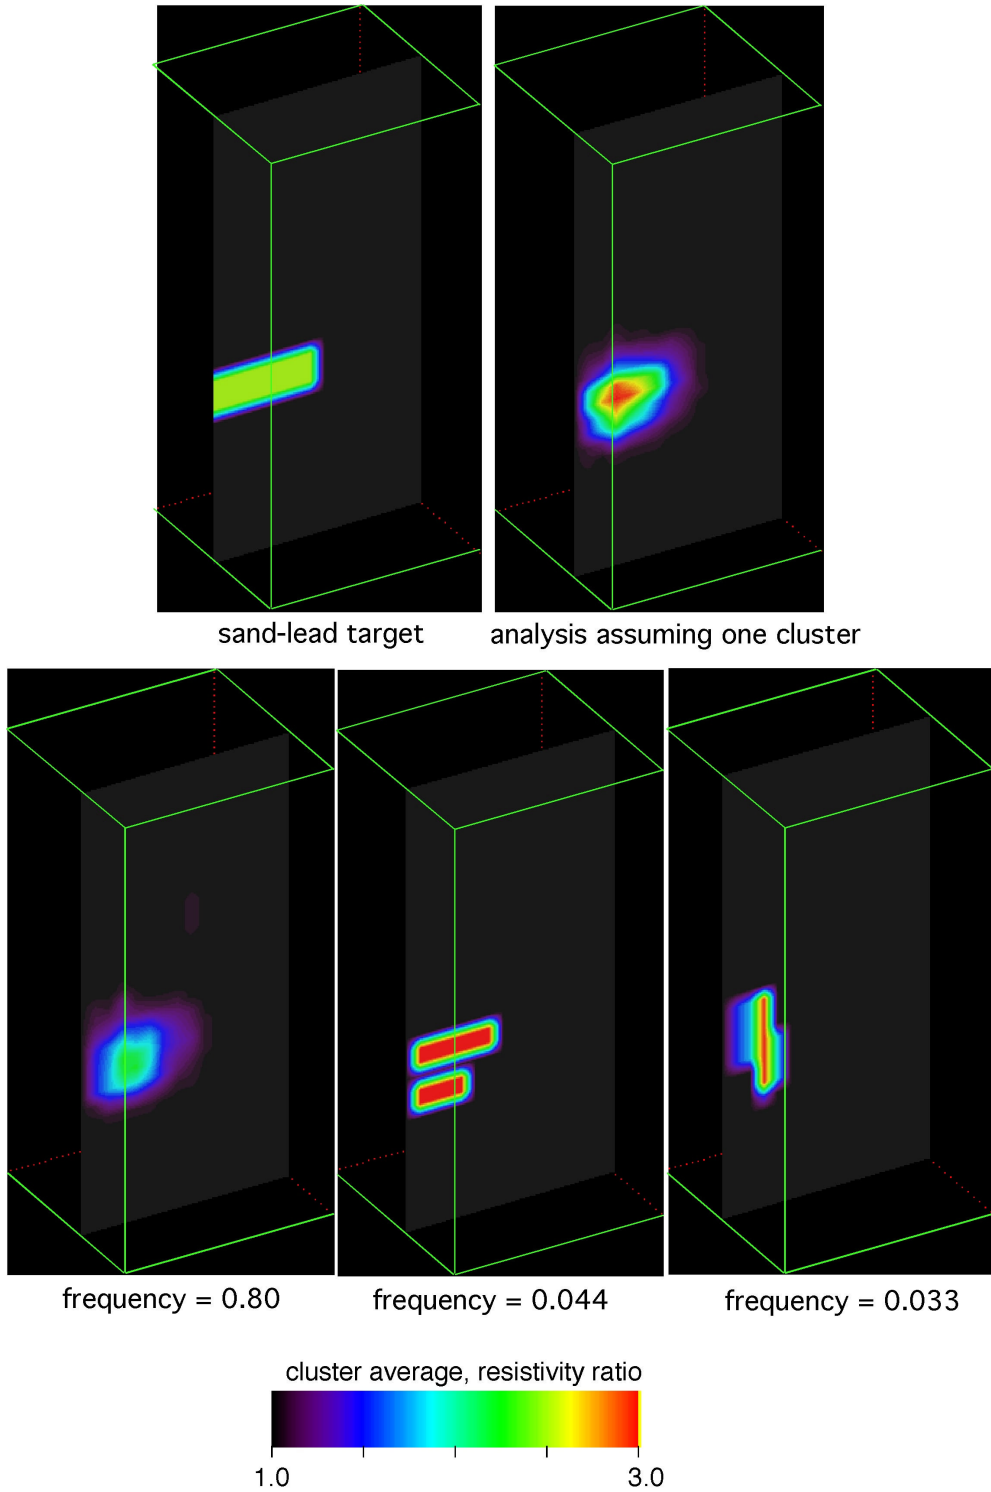


Figure 10 presents clustering analysis results for the case of the sand-lead target. The top left frame shows a vertical slice through the target. The right frame shows the voxel-wise average resistivity obtained when a uni-modal posterior distribution is assumed. The bottom frames shows the average resistivity ratio obtained when a multi-modal distribution is assumed. The three most probable clusters and their corresponding frequencies are shown.

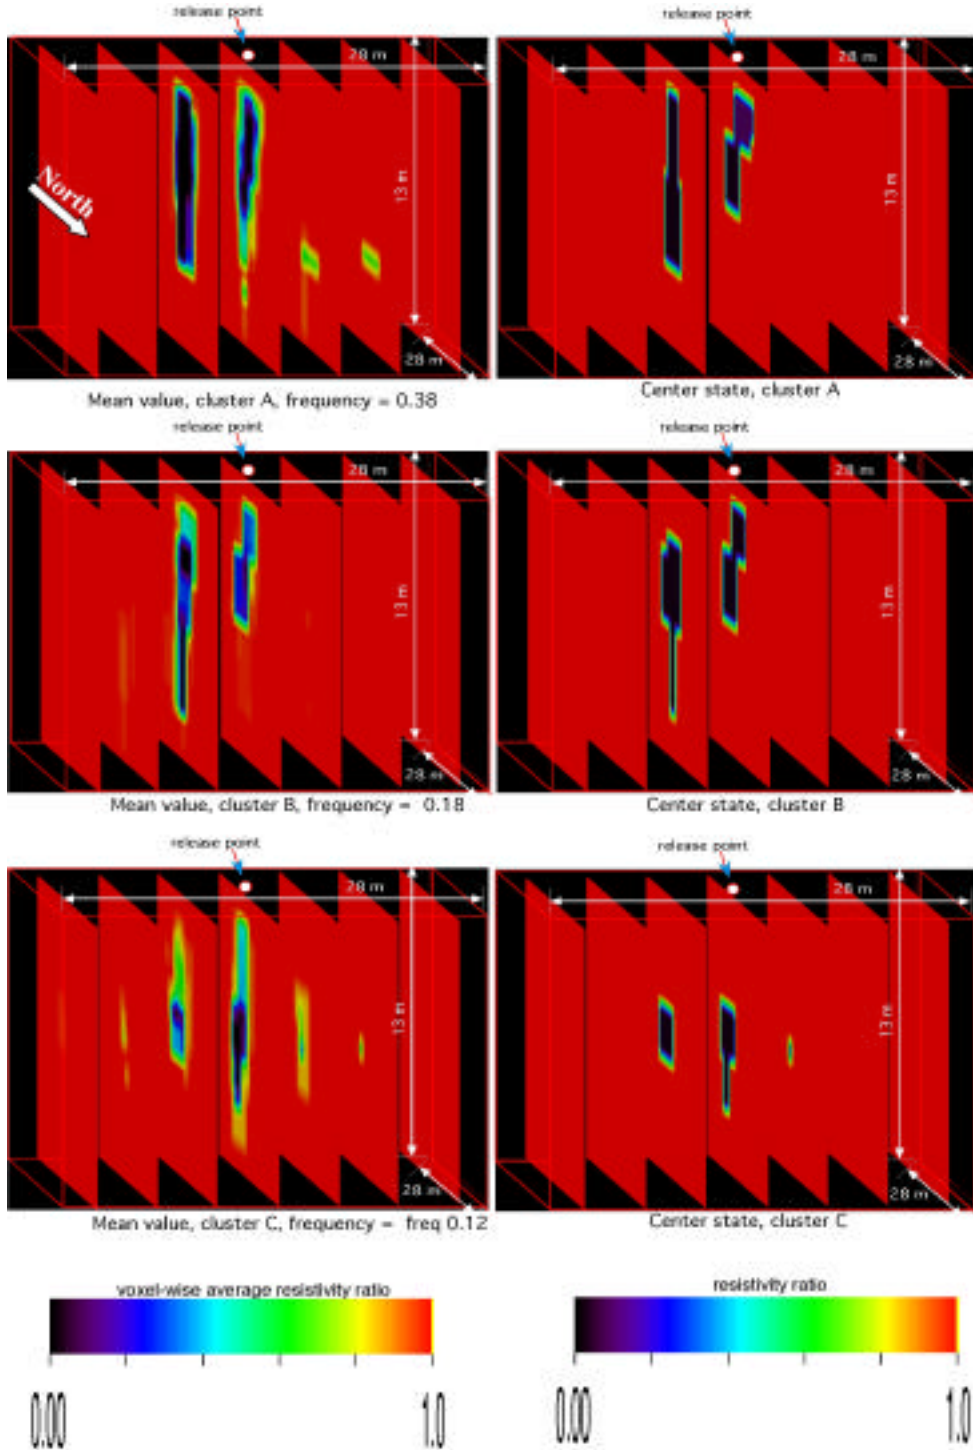


Figure 11 displays clustering analysis results corresponding to the tank release experiment. The left column of frames show the voxel-wise average resistivity ratio for the top three most probable clusters. The right column of frames shows the center state for the three most probable clusters.

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551
